

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Method and Apparatus for Creating Templates

Inventors:

Lev Novik

Patrick R. Kenny

Alexander E. Nosov

ATTORNEY'S DOCKET NO. MS1-702US

1 **RELATED APPLICATIONS**

2 This application is a continuation-in-part of co-pending application Ser. No.
3 09/847,534, filed May 1, 2001, entitled "Method and Apparatus for Correlating
4 Events", and incorporated herein by reference.

5

6 **TECHNICAL FIELD**

7 The present invention relates to computing systems and, more particularly,
8 to a language for describing templates used to collect and correlate various events
9 generated throughout a computing environment.

10

11 **BACKGROUND**

12 Computer systems, such as servers and desktop personal computers, are
13 expected to operate without constant monitoring. These computer systems
14 typically perform various tasks without the user's knowledge. When performing
15 these tasks, the computer system often encounters events that require a particular
16 action (such as logging the event, generating an alert for a particular system or
17 application, or performing an action in response to the event). Various
18 mechanisms are available to handle these events.

19 A computing enterprise typically includes one or more networks, services,
20 and systems that exchange data and other information with one another. The
21 enterprise may include one or more security mechanisms to safeguard data and
22 authenticate users and may utilize one or more different data transmission
23 protocols. At any particular time, one or more networks, services or systems may
24 be down (e.g., powered down or disconnected from one or more networks).

1 Networks, services or systems can be down for scheduled maintenance, upgrades,
2 overload or failure. Application programs attempting to obtain event data must
3 contend with the various networks, services, and systems in the enterprise when
4 they are down. Additionally, application programs must contend with the security
5 and network topology limitations of the enterprise as well as the various protocols
6 used in the enterprise.

7 Operating system components, services, and applications generate a variety
8 of different events. A particular component or application may request to be
9 informed of a particular event (e.g., when a server crashes or when a user logs on
10 to the system). Other components or applications may want to be notified when a
11 particular series of events occur within a particular time period. For example, a
12 network administrator may want to know when a server crashes within three
13 seconds of a user logging into the system. Server crashes alone may be relatively
14 common and user logins may also be common such that the network administrator
15 is not particularly interested in either event by itself. However, when these two
16 events occur within a few seconds of one another, there may be a relationship
17 between the two events (e.g., the user login was at least partially responsible for
18 the server crash).

19 Existing systems provide predefined functions that allow a network
20 administrator or other user to create a relationship between two events. This
21 relationship between two events is commonly referred to as a “correlation”
22 between the two events. The predefined correlation functions provided by existing
23 systems require the user to select from one of the predefined functions. If the
24 correlation function desired by the user has not already been created, the user must
25

1 request that the developer or supplier of the functions create a new function to
2 meet the user's needs. If the developer or supplier is willing to create a new
3 correlation function, this custom development work may be very expensive.
4 Depending on the expected demand for the new correlation function, the developer
5 or supplier may not be willing to create the requested function.

6 If the developer is unwilling to create a new correlation function or the cost
7 is too high, the user can attempt to use an existing correlation function that is
8 "closest" to the user's requirements. Such a solution may result in a significant
9 number of unwanted event notifications or may result in a failure to notify the user
10 of a desired sequence of events.

11

12 **SUMMARY**

13 The system and method described herein addresses the limitations of the
14 prior art by providing a template description language that allows a user to
15 develop functions that identify certain events and correlations between multiple
16 events and/or data.

17 The system and method described herein supports a template description
18 language that allows a user to create templates that are used to correlate multiple
19 events and/or data elements. The templates created using the template description
20 language in turn create state machines to execute a function, such as an event
21 correlation function. The template description language supports the development
22 of two general types of templates: correlation helper templates and correlation
23 scenario templates. A correlation helper template is used to parameterize a
24 particular aspect of a correlation scenario. The correlation helper template does

1 not define a complete correlation scenario, but instead leaves one or more
2 parameters to be defined by the user implementing the template. A correlation
3 scenario template defines a specific correlation scenario and maintains its own
4 correlation state. Correlation scenario templates can be used as building blocks to
5 create larger, more complex correlation scenario templates. The templates
6 described herein allow a user to develop templates to perform any desired event
7 and/or data correlation functions.

8 In one embodiment, a template creator creates a template class that stores at
9 least one template parameter. The template creator specifies at least one
10 additional object to be created when an instance of the template class is created. A
11 template user initiates the creation of a new instance of the template class. If an
12 error occurs during the creation of the new instance of the template class, the
13 template user receives a report identifying the error.

14 Another embodiment creates a first template class and designates inputs
15 and outputs associated with the first template class. A second template class is
16 created and inputs and outputs associated with the second template class are
17 designated. A third template class is created by combining the first template class
18 and the second template class.

19 In a particular embodiment, a template description structure includes at
20 least one template class that stores multiple template parameters. The template
21 description structure also includes at least one template builder class that identifies
22 at least one additional object to be created with each instance of the template class.
23 The template description structure further includes at least one order parameter
24 that identifies the order in which the template builder classes are instantiated.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a block diagram of a system that receives event information from multiple event providers and provides event information to multiple event consumers.

Fig. 2 is a flow diagram illustrating an event-handling procedure.

Fig. 3 illustrates a block diagram of a system that receives multiple events from various event sources, filters the events, and distributes the filtered events to multiple event consumers.

Fig. 4 is a flow diagram illustrating a procedure for handling events in the system of Fig. 3.

Fig. 5 illustrates a block diagram of an exemplary system in which a correlator receives events, data, and correlation functions from multiple sources.

Fig. 6 is a flow diagram illustrating a procedure for correlating events and/or data in an enterprise.

Fig. 7 is a flow diagram illustrating a procedure for implementing a state machine that applies a correlation function.

Fig. 8 is a flow diagram illustrating a procedure for defining a template class and creating a new instance of the template class.

Fig. 9 is a flow diagram illustrating a procedure for combining two templates to create a third template.

Fig. 10 illustrates an example of a suitable operating environment in which the event distribution and event handling system and method may be implemented.

DETAILED DESCRIPTION

The system and method described herein provides a template description language that allows a user to define templates that perform various correlation functions between events and/or data. A correlation helper template is used to parameterize one or more aspects of a correlation scenario. The correlation helper template leaves one or more parameters to be defined by the user creating an instance of the template. A correlation scenario template may be used as a building block to create larger, more complex correlation scenario templates.

Web-Based Enterprise Management (WBEM) provides uniform access to management information throughout an enterprise. WBEM is an industry initiative to develop technology for accessing management information in an enterprise environment. This management information includes, for example, information on the state of system memory, inventories of currently installed client applications, and other information related to the status of the system. A particular embodiment of the event-handling system is represented using Windows Management Instrumentation (WMI) developed by Microsoft Corporation of Redmond, Washington, which provides an infrastructure to handle various events generated by event sources throughout an enterprise.

WMI technology enables systems, applications, networks, and other managed components to be represented using the Common Information Model (CIM) designed by the Distributed Management Task Force (DMTF). This model is used to perform correlation functions discussed herein. CIM is an extensible data model for representing objects that exist in typical management environments. CIM is able to model anything in the managed environment,

1 regardless of the location of the data source. The Managed Object Format (MOF)
2 language is used to define and store modeled data. In addition to data modeling,
3 WMI provides a set of base services that include query-based information retrieval
4 and event notification. Access to these services and to the management data is
5 provided through a single programming interface.

6 WMI classes define the basic units of management. Each WMI class is a
7 template for a type of managed object. For example, Win32_DiskDrive is a model
8 representing a physical disk drive. For each physical disk drive that exists, there is
9 an instance of the Win32_DiskDrive class. WMI classes may contain properties,
10 which describe the data of the class and methods, which describe the behavior of
11 the class.

12 WMI classes describe managed objects that are independent of a particular
13 implementation or technology. WMI includes an eventing subsystem that follows
14 the publish-subscribe model, in which an event consumer subscribes for a
15 selection of events (generated by one or more event providers) and performs an
16 action as a result of receiving the event. WMI also provides a centralized
17 mechanism for collecting and storing event data. This stored event data is
18 accessible by other systems via WMI tools and/or application programming
19 interfaces (APIs).

20 Although particular embodiments are discussed herein as using WMI,
21 alternate embodiments may utilize any enterprise management system or
22 application, whether web-based or otherwise. The event providers and event
23 consumers discussed herein are selected for purposes of explanation. The
24 teachings of the present invention can be used with any type of event provider and
25

any type of event consumer. Additionally, the event-handling system and method described herein can be applied to any type of enterprise or other arrangement of computing devices, applications, and/or networks.

Fig. 1 illustrates a block diagram of a system 100 that receives event information from multiple event providers 114 (i.e., event sources) and provides event information to multiple event consumers 102 (i.e., the users of the event data). System 100 includes a WMI module 106, which receives event data from multiple event sources 114 and receives requests for information (e.g., notification of particular events) from multiple event consumers 102. The multiple event sources are identified as event providers 112. The multiple event consumers are identified as applications 104.

Event providers 112 include, for example, systems, services or applications that generate event data. An exemplary event provider is a disk drive (or an application that monitors the status of a disk drive). The disk drive may generate an event indicating the available storage capacity on the disk drive or indicating the amount of data currently stored on the disk drive. The disk drive may also generate an event indicating that the disk drive is nearly full of data (e.g., when ninety-five percent or more of the disk drive's capacity is used).

Event consumers 102 may request to be notified of certain events (also referred to as “subscribing” to an event). An example event consumer is an application that manages multiple storage devices in an enterprise. The application may request to receive events generated by any of the disk drives or other storage devices in the enterprise. The application can use this event information to distribute storage tasks among the multiple storage devices based

1 on the available capacity of each device and/or the quantity of read or write
2 requests received by each storage device.

3 System 100 also includes a set of policies 110, which are accessible by
4 WMI module 106. Policies 110 may control the configuration of one or more
5 systems in the enterprise. Other policies may define various activities, such as
6 event filtering, event correlation, and the forwarding of events to particular
7 devices or applications. A database 108 is coupled to WMI module 106.
8 Database 108 stores various information related to the enterprise. For example,
9 database 108 can store event data (i.e., creating an event log), policy data, and
10 enterprise configuration information.

11 The WMI module 106 uses WMI features to provide a distributed
12 architecture that is capable of selecting, filtering, correlating, forwarding, storing,
13 and delivering event data in an enterprise. The WMI module also allows event
14 consumers to request data related to a particular event, request data from a
15 particular node or device in the enterprise, define the manner in which events are
16 correlated with one another, define how certain events should be forwarded, and
17 define how to store event data.

18 The WMI module 106 provides a policy-based administration of the
19 enterprise. The policy infrastructure allows administrators to set a policy in the
20 Directory Service (DS) and the WMI module ensures that the proper set of WMI
21 objects (e.g., filters, bindings, correlators, consumers, and configuration objects)
22 are delivered to the proper devices or applications in the enterprise.

23 As shown in Fig 1, policies 110 and database 108 are separate from WMI
24 module 106. However, in alternate embodiments, policies 110 and/or database
25 108 may be integrated into WMI module 106.

1 Table 1 below identifies various types of event providers available in a
2 particular embodiment. Additionally, the table includes a description of the events
3 generated by each event provider. For example, the Win32 Provider generates
4 events that include information related to the operating system, computer system,
5 peripheral devices, file systems, and security for a particular device (such as a
6 computer system) in the enterprise.

7
8 TABLE 1
9

10 Event Provider	11 Description of Events Provided
12 Win32 Provider	13 Supplies information about the 14 operating system, computer system, 15 peripheral devices, file systems, and 16 security.
17 WDM Provider	18 Supplies low-level Windows Driver 19 Model (WDM) information for user 20 input devices, storage devices, network 21 interfaces, and communications ports.
22 Performance 23 Counter Provider	24 Exposes the raw performance counter 25 information used to compute various 26 performance values.
27 Windows Installer 28 Provider	29 Supplies information about 30 applications installed with the 31 Windows Installer.

32
33 Fig. 2 is a flow diagram illustrating an event-handling procedure 200. The
34 WMI module monitors event activity throughout the enterprise (block 202). The
35 procedure 200 determines whether event data has been received from an event
36 provider (block 204). If event data has been received, the WMI module records
37 the event data and initiates any appropriate actions (block 206). An example
38 action includes notifying an event consumer of the event (e.g., if the event
39 consumer previously subscribed to such an event).

At block 208, the procedure 200 determines whether a new subscription for event information has been received. The procedure 200 may also determine whether a request to revise an existing subscription has been received. If a new subscription (or a revised subscription) is received, the procedure continues to block 210 where the WMI module retrieves the requested event information and provides the information to the requesting event customer. Alternatively, the procedure may log the subscription request and notify the requesting event consumer when the next event is received that qualifies under the consumer's subscription request.

Fig. 3 illustrates a block diagram of a system 300 that receives multiple events from various event sources, filters the events, and distributes the filtered events to multiple event consumers. Multiple events 302 are received by an event filter 304 that determines which of the received events are passed through the filter to a correlator 310. The filter 304 applies various filter criteria 308 to the received events in determining which events pass through the filter to correlator 310. Additional details regarding correlator 310 and the correlation functions are discussed below.

The correlator 310 correlates various events and creates additional events 312 that are provided to multiple filters 314, 320, 326, and 332. Each filter 314, 320, 326, and 332 includes various filter criteria that determines what event characteristics are required to allow the event to pass through the filter. Although each event 312 is sent to all four filters, the event may be rejected (i.e., not pass through the filter) by any or all of the filters. Similarly, a particular event may pass through two or more different filters, depending on the filter criteria associated with each filter.

1 Each filter 314, 320, 326, and 332 is associated with a consumer (i.e., an
2 event consumer) 316, 322, 328, and 334, respectively. For example, events that
3 pass through filter 314 are provided to event logging consumer 316, which logs
4 the event data to a storage device 318. The logged data can be retrieved at a later
5 time for analysis or other purposes. Events that meet the criteria of filter 320 are
6 provided to event forwarding consumer 322, which generates a forwarded event
7 324 that is distributed to one or more destinations. Events that satisfy the criteria
8 of filter 326 are provided to mail consumer 328, which generates and sends an
9 email message 330 in response to receipt of each event. The email message 330
10 may contain information about one or more events (such as the event type or the
11 source of the event). Events that pass through filter 332 are provided to scripting
12 consumer 334, which executes a script that may perform a function and/or
13 generate a script output 336.

14 Although the example of Fig. 3 illustrates four filters 314, 320, 326, and
15 332 (and associated consumers 316, 322, 328, and 334, respectively) that receive
16 events 312, alternate embodiments may include any number of filters and
17 associated consumers. Further, the actions performed by consumers 316, 322,
18 328, and 334 are provided as examples. Alternate consumers may perform any
19 type of action in response to receiving an event.

20 Fig. 4 is a flow diagram illustrating a procedure 400 for handling events in
21 the system of Fig. 3. An event is received by an event filter (block 402), such as
22 filter 304 in Fig. 3. The procedure 400 determines whether the received event
23 satisfies the event filter (block 404). Satisfying the event filter includes satisfying
24 the filter criteria (e.g., filter criteria 308). If the received event does not satisfy the
25 event filter, then the received event is discarded (block 406). Discarding an event

1 may include ignoring the event or deleting the event and any reference to the event
2 from storage registers or other storage mechanisms. If the received event satisfies
3 the event filter (i.e., passes through the filter), a correlator correlates multiple
4 received events (block 408) and may generate one or more new events that are
5 provided to multiple event filters (e.g., filters 314, 320, 326, and 332 in Fig. 3) in
6 block 410.

7 Each event filter analyzes the event using its own filter criteria (block 412).
8 Next, each event filter determines whether the event meets the event filter's
9 criteria (block 414). This determination is performed by each event filter based on
10 the filter criteria for that particular event filter. If the event does not meet the
11 criteria for a particular event filter, that event filter discards the event (block 416).
12 However, if the event satisfies the criteria for a particular event filter, that event
13 filter passes the event to the event consumer that corresponds to the particular
14 event filter (block 418). The event consumer then performs one or more actions
15 based on the event (block 420). For example, the actions may include generating
16 an email message or forwarding the event to another system. The procedure of
17 Fig. 4 is repeated for each received event.

18 Fig. 5 illustrates a block diagram of an exemplary system 500 in which a
19 correlator 502 receives events 504, data elements 506, and correlation functions
20 508 from multiple sources. Events 504 have passed through a filter (e.g., filter
21 304 in Fig. 3) prior to being received by correlator 502. The terms "data" and
22 "data elements" are used interchangeably herein. Correlator 502 may be similar to
23 the correlator 310 (Fig. 3) discussed above. Correlator 502 may receive events
24 504 directly from an event source or may receive events through an event filter,
25 such as filter 304 discussed above. In one embodiment, correlator 502 receives

1 specific event information based on event subscriptions applied to various event
2 sources in the enterprise. Correlator 502 may receive data elements 506 directly
3 from a data source or may receive the data through an intermediate device, such as
4 a database or other data storage or data collection device. Correlator 502 also
5 receives correlation functions from one or more sources, such as from an
6 administrator's node. The correlation function 508 is a group of updating
7 consumers (discussed below) that receive events like other consumers. These
8 updating consumers cause one or more events 510 to occur (either directly or
9 indirectly).

10 Correlator 502 applies the received correlation functions to the events and
11 data received from various sources throughout the enterprise. When the
12 conditions of a particular correlation function are satisfied, correlator 502
13 generates an event 510, which is distributed to various event consumers in the
14 enterprise. In one embodiment, the event 510 is provided to event consumers that
15 subscribed to receive that particular event.

16 Fig. 6 is a flow diagram illustrating a procedure 600 for correlating events
17 and/or data in an enterprise. A correlator (e.g., correlator 502 in Fig. 5) receives
18 events, data and correlation functions from multiple sources in an enterprise (block
19 602). The received events and data are compared to the correlation functions
20 (block 604). The procedure 600 next determines whether any of the correlation
21 functions are satisfied by the received events and data (block 606). If not, the
22 procedure 600 continues receiving events, data and correlation functions (block
23 608), and returns to block 604 to continue comparing received events and data to
24 the correlation functions. If any of the correlation functions are satisfied by the
25 received events and data, the correlator generates one or more events based on the

1 satisfied correlation function (block 610). After generating one or more events in
2 block 610, the procedure 600 continues to block 608 to continue receiving events,
3 data and correlation functions.

4 Fig. 7 is a flow diagram illustrating a procedure 700 for implementing a
5 state machine that is capable of implementing a correlation function. Initially, a
6 user defines or selects a pre-defined schema for the desired state machine (block
7 702). Next, an instance of the desired state machine is created (block 704).
8 Transitions for the state machine are defined by subscribing to one or more events
9 (block 706). An update consumer is applied to the state machine to update the
10 state of the state machine (block 708). The updating consumer is applied to the
11 state machine each time an event to which the updating consumer has subscribed
12 occurs.

13 Procedure 700 then determines whether the state machine is in its final state
14 (block 710). If the state machine is not in its final state, the procedure 700 returns
15 to block 708 to again apply the update consumer to the state machine. If the state
16 machine is in its final state, the procedure 700 continues to block 712, which
17 deletes the current instance of the state machine. If another correlation function is
18 to be implemented, a new instance of the desired state machine is created and
19 executed.

20 Examples of events include a server crash, a user logging into the system,
21 or a particular device becoming unavailable. Example data elements include the
22 available disk space, the current memory utilization, and the number of users
23 logged into particular servers. An example correlation function that correlates two
24 events generates an email message when two different server crashes occur within
25 five second of one another. An example correlation function that correlates an

1 event with data generates an event when a server crashes and the available storage
2 space on the server's hard drive is less than five megabytes. Another example
3 correlation function pages an administrator when the available storage space on a
4 server's hard disk stays below ten megabytes for at least five minutes. Any other
5 selection of events and/or data can be combined to create a correlation function
6 based on the desires of an administrator or other user.

7 As mentioned above, a state machine implements a desired correlation
8 function that correlates events and/or data. A set of commonly used state
9 machines are provided for use by administrators (or other users) in defining
10 correlation functions. These commonly used state machines require the
11 administrator to fill in certain parameters, but the administrator is not required to
12 understand the programming language used to create the state machine. If the set
13 of commonly used state machines does not include a state machine that performs
14 the desired correlation function, a new state machine can be created using the
15 appropriate programming language. The programming language can be any
16 database language or other non-procedural language. In a particular embodiment,
17 the programming language is SQL.

18 Each state machine is a class object. One or more instances of a state
19 machine can be implemented simultaneously to monitor different events and data.
20 In a particular enterprise, any number of instances of state machines may be
21 operating simultaneously. In one embodiment of the invention, SQL is used to
22 query various states in any state machine.

23 In a particular example, the schema for a state machine that detects a
24 specific number of process crashes within a specified time period can be defined
25 as follows.

```
1 Class StateA
2 {
3     string ProcessName;
4     int NumCrashes;
5     int RemainingTime;
6 }
```

In the above example, the state machine is a class having three properties. ProcessName is a string that represents the name of the process being monitored for crashes. NumCrashes is an integer that identifies the number of crashes that will trigger an event if those crashes occur within a particular time period, which is defined by the property RemainingTime, which is an integer value. The RemainingTime property is reset each time a new instance of StateA is created. If RemainingTime reaches zero without being reset, an event is triggered indicating that the state machine time expired before detecting the specified number of crashes. When RemainingTime reaches zero, that particular instance of StateA is deleted because the specified parameters were not satisfied within the particular time period.

An administrator wanting to use the correlation function defined by StateA first creates an instance of StateA. The administrator then provides a value for NumCrashes and RemainingTime. Thus, the administrator need not understand the complete syntax of the state machine and need not understand the programming language used to define and create the state machine.

After defining the schema for the StateA state machine, the transitions for the state machine (i.e., the transitions from one state to another) are defined by subscribing to various events. Specifically, the transition is defined by the

1 updating consumer and the event that causes the transition is defined by the event
2 subscription. These event subscriptions function as the transitions for the state
3 machine. When an appropriate event occurs, the state machine transitions to the
4 next state. The state machine transitions are defined by identifying the event that
5 will cause the transition and identifying the action to perform based on the
6 occurrence of the event. The action may include, for example, generating an email
7 message, logging event data to a file, or forwarding an event to one or more
8 destinations. The transitions are defined using updating consumer instances.

9 After defining the transitions for the state machine, an updating consumer is
10 used to update the state of the state machine. The updating consumer (named
11 “update”) is a class object. One or more instances of the updating consumer can
12 be implemented simultaneously to handle the updating of different state machines.
13 An example updating consumer implementation is illustrated below.

14 Update StateA where ProcessName = ThisEvent.ProcessName
15 set NumCrashes = NumCrashes + 1

16
17 In this example, the updating consumer updates an instance of state machine
18 StateA, defined above. The ProcessName property is defined as
19 “ThisEvent.ProcessName”, which inserts the name of the process that crashed
20 (which is identified in the received crash event) as “ThisEvent”. The property
21 NumCrashes is incremented by one each time a crash event is received.

22 While a particular state machine is operating, the various internal states of
23 the state machine can be obtained (e.g., queried). This allows an administrator or
24 other user to observe the correlation as the various events occur in a system. Even
25 if the conditions have not yet been met to generate the appropriate event, the

1 administrator can observe the current state or value of different properties (e.g.,
2 how many crashes have occurred or how much time is left before the state
3 machine is reset). The ability to observe the various states and properties of the
4 state machine assists with troubleshooting and determining whether the desired
5 correlation function has been properly established.

6 Various examples have been discussed herein in which two different events
7 are correlated with one another or an event is correlated with data. However, in
8 alternate embodiments, any number of events can be combined together to form a
9 correlation function. Similarly, any number of events can be combined with one
10 or more data elements to create a correlation function.

11 The following example illustrates classes and class instances, a correlation
12 scenario, updating consumers, filters and bindings as used with the present
13 invention. Example class and instances of the class:

```
14 class ExampleClass
15 {
16     [key] string Name;
17     boolean Prop;
18 };
19
20 instance of ExampleClass
21 {
22     Name = "A";
23     Prop = TRUE;
24 };
25
26 instance of ExampleClass
27 {
28     Name = "B";
29     Prop = FALSE;
30 };
31
32
33
```

25 The correlation scenario is defined:

```

[ dynamic, provider("Microsoft WMI Transient Provider")]
class ExampleCorrelationState : MSFT_CorrelationStateBase
{
    boolean ReceivedEventA;
    boolean ReceivedEventB;
    [trns_egg_timer] uint32 Timer;
};

Class BothAandBEvent : MSFT_UCEventBase
{
    string Name;
};

instance of MSFT_UpdatingConsumer as $UI
{
    Id = "Initializer";
    Scenario = "ExampleCorrelationScenario";
    Commands = { "INSERT INTO ExampleCorrelationState "
        "( Id, Scenario, ReceivedEventA, ReceivedEventB, Timer ) "
        " ( 'ExampleCorrelationState', 'ExampleCorrelationScenario', "
        " FALSE, FALSE, 0 )"};
};

instance of MSFT_UpdatingConsumer as $UA
{
    Id = "SetEventA";
    Scenario = "ExampleCorrelationScenario";
    Commands = { "UPDATE ExampleCorrelationState "
        "SET ReceivedEventA = TRUE, Timer = 5 "
        "WHERE Scenario = 'ExampleCorrelationScenario' "};
};

instance of MSFT_UpdatingConsumer as $UB
{
    Id = "SetEventB";
    Scenario = "ExampleCorrelationScenario";
    Commands = { "UPDATE ExampleCorrelationState "
        "SET ReceivedEventB = TRUE, Timer = 5 "
        "WHERE Scenario = 'ExampleCorrelationScenario' "};
};

```

1 The \$UA and \$UB updating consumers cause the timer to be reset to five seconds
2 whenever either EventA or Event B occurs. The next updating consumer causes
3 the ReceivedEventA and the ReceivedEventB to be reset when the timer expires.

```
4 instance of MSFT_UpdatingConsumer as $UTE
5 {
6     Id = "ResetTimer";
7     Scenario = "ExampleCorrelationScenario";
8     Commands = { "UPDATE ExampleCorrelationState "
9                 "SET ReceivedEventA = FALSE, ReceivedEventB = FALSE "
10                "WHERE Scenario = 'ExampleCorrelationScenario' "};
11    };
12
```

13 The following defines filters and bindings to fully define the scenario:

```
14 instance of __EventFilter as $FSC
15 {
16     Name = "ScenarioCreation";
17     Query = "SELECT * FROM __InstanceCreationEvent "
18         "WHERE TargetInstance ISA 'MSFT_UCScenario' "
19         "AND TargetInstance.Id = 'ExampleCorrelationScenario'";
20     QueryLanguage = "WQL";
21 };
22
23 instance of __EventFilter as $FSM
24 {
25     Name = "ScenarioModification";
26     Query = "SELECT * FROM __InstanceModificationEvent "
27         "WHERE TargetInstance ISA 'MSFT_UCScenario' "
28         "AND TargetInstance.Id = 'ExampleCorrelationScenario'";
29     QueryLanguage = "WQL";
30 };
31
32 instance of __EventFilter as $FBOOT
33 {
34     Name = "OnBoot";
35     Query = "SELECT * FROM MSFT_TransientRebootEvent ";
36     QueryLanguage = "WQL";
37 };
38
39 instance of __EventFilter as $FA
40 {
41 }
```

```

1   Name = "EventAFilter";
2   Query = "SELECT * FROM __InstanceModificationEvent "
3     "WHERE TargetInstance ISA 'ExampleClass' "
4     "AND TargetInstance.Name = 'A'";
5   QueryLanguage = "WQL";
6   };
7
8   instance of __EventFilter as $FB
9   {
10    Name = "EventBFilter";
11   Query = "SELECT * FROM __InstanceModificationEvent "
12     "WHERE TargetInstance ISA 'ExampleClass' "
13     "AND TargetInstance.Name = 'B'";
14   QueryLanguage = "WQL";
15   };
16
17   instance of __EventFilter as $FTE
18   {
19    Name = "TimerExpiredEvent";
20    Query = "SELECT * FROM MSFT_TransientEggTimerEvent "
21      "WHERE Object ISA \"ExampleCorrelationState\" "
22      "AND Object.Scenario = 'ExampleCorrelationScenario'";
23    QueryLanguage = "WQL";
24   };
25

```

Defining the bindings:

```

16   instance of __FilterToConsumerBinding
17   {
18    Filter = $FSC;
19    Consumer = $UI;
20   };
21
22   instance of __FilterToConsumerBinding
23   {
24    Filter = $FSM;
25    Consumer = $UI;
26   };
27
28   instance of __FilterToConsumerBinding
29   {
30    Filter = $FBOOT;
31    Consumer = $UI;
32   };
33
34
35

```

```
1     };
2
3     instance of __FilterToConsumerBinding
4     {
5         Filter = $FA;
6         Consumer = $UA;
7     };
8
9     instance of __FilterToConsumerBinding
10    {
11        Filter = $FB;
12        Consumer = $UB;
13    };
14
15    instance of __FilterToConsumerBinding
16    {
17        Filter = $FTE;
18        Consumer = $UTE;
19    };
20
```

When creating this updating consumer scenario, the activation can be triggered to occur using the MSFT_UCScenario instance. Since, in this example, the system would have \$CI consumer tied to the creation of the Scenario instance, the following instantiation would cause the initialization to occur:

```
16     instance of MSFT_UCScenario
17     {
18         Id = "ExampleCorrelationScenario";
19         Name = "ExampleCorrelationScenario";
20     };
21
```

This instance helps the updating consumer provider determine how state instances relate to the scenario:

```
22     instance of MSFT_UCScenarioAssociationInfo
23     {
24         Id = "StateAssociation";
25         Scenario = "ExampleCorrelationScenario";
26         Query = "SELECT * FROM ExampleCorrelationState "
```

```
1      "WHERE Scenario = 'ExampleCorrelationScenario'";
2  };
```

2 Finally, an example filter to determine when both events occurred within the
3 windows:

```
4      instance of __EventFilter as $FBOTH
5  {
6      Name = "BothEventsOccurred";
7      Query = "SELECT * FROM __InstanceModificationEvent "
8          "WHERE TargetInstance ISA \"ExampleCorrelationState\" "
9          "AND TargetInstance.ReceivedEventA = TRUE "
10         "AND TargetInstance.ReceivedEventB = TRUE "
11         "AND TargetInstance.Scenario = 'ExampleCorrelationScenario' ";
12     QueryLanguage = "WQL";
13 };
```

11 To create a custom event that is triggered when this condition is met, then the user
12 can subscribe an updating consumer to it:

```
13      instance of MSFT_UpdatingConsumer as $UBOTH
14  {
15      Id = "BothEventsOccurred";
16      Scenario = "ExampleCorrelationScenario";
17      Commands = { "INSERT INTO BothAandBEvent ( Name ) "
18          "( 'ExampleCorrelationScenario' ) " };
19  };
20
21
22      instance of __FilterToConsumerBinding
23  {
24      Filter = $FBOTH;
25      Consumer = $UBOTH;
26  };
27
28
29
30
31
32
33
34
35
```

Templates

23 The systems discussed above provide various state machines and other
24 tools to correlate events and/or data. The use of templates, discussed below,
25

1 provide another tool that allows a user to create a parameterized scenario for the
2 benefit of many other users. Once a template is created, many users can utilize the
3 template by simply specifying the appropriate parameter values. As used herein, a
4 “template creator” is an individual that creates one or more templates that are used
5 by one or more “template users”. The templates created using the template
6 description language described herein can be used to create a state machine for a
7 correlation scenario.

8

9 Template Provider

10 A “template provider” is at the core of the implementation of templates. In
11 a particular embodiment, the template provider is a WMI Instance Provider. The
12 template provider was developed to help administrators handle complex updating
13 of consumer scenarios.

14 Each time a template user creates an instance of a template class (discussed
15 below), the template provider creates a group of other instances, referred to as
16 “target objects”, which are specified by the template creator at the time the
17 template is created. The values in the target objects are parameterized by the
18 values of the instance of the template class (these values are provided by the
19 template user). This process is referred to as instantiation. When an instance of a
20 template is deleted, the template provider automatically deletes all of the target
21 objects that were created when the template instance was created.

22 Fig. 8 is a flow diagram illustrating a procedure 800 for defining a template
23 class and creating a new instance of the template class. Initially, a template
24 creator creates a template class that will store template parameters (block 802).
25 The template creator then specifies additional objects that need to be created when

1 an instance of the template class is created (block 804). After the template class
2 has been defined (i.e., the two steps above), a new instance of the template class
3 can be created by a template user (block 806). The template user creates the new
4 instance of the template class by providing one or more parameters needed by the
5 template class.

6 If a failure occurs during the instantiation of the template class, the error is
7 reported to the template user (block 812). If the template instance is created
8 without error, the template user is then able to utilize the new instance of the
9 template class (block 810). Additionally, other users may utilize the new instance
10 of the template class or may create another instance of the same template class
11 using the procedures discussed above.

12

13 Template Class

14 As mentioned above with respect to Fig. 8, the first step in creating a
15 template is the creation of a class that will store template parameters. This class is
16 referred to as a “template class” and its properties are referred to as template
17 arguments. An instance of the template class will be referred to as a template
18 instance. The template class is created by the designer of the template, and is
19 designated as provided by the template provider (described below). For instance,
20 the following is a valid template class:

21 [dynamic, provider("Microsoft WMI TemplateProvider")]
22 Class CorrelationTemplate : MSFT_TemplateBase
23 {
24 [key] string Scenario;
25 string WatchClass;
26 string WatchScopeExpr;
27 string WatchProp;

1 sint32 SomeValue;
2 };

3 In the above example, the class is derived from MSFT_TemplateBase. The use of
4 MSFT_TemplateBase is optional. The schema of MSFT_TemplateBase is:

5 [abstract]
6 class MSFT_TemplateBase
7 {
8 [NotNull] boolean Active = TRUE;
9 };

10 The “Active” property allows a template instance to exist, but not the
11 objects that the template instantiates. This is useful for storing template instances
12 that will be activated at a later time.

13

14 Template Instantiation

15 Template instantiation can occur when a template instance is created or
16 modified. Template instantiation will take place if the template class does not
17 inherit from MSFT_TemplateBase or the template class does inherit from
18 MSFT_TemplateBase and the active property is “True”. When a template
19 instance is modified, (e.g., one or more of its template arguments have changed),
20 then only target objects that change will be re-instantiated. If the template
21 instance is deleted, then all target objects that it owns will be deleted as well.

22

23 Template Builders

24 The next step in creating a template is specifying what lower-level objects
25 need to be created when a user creates an instance of the template. This is done by

1 creating instances of the “TemplateBuilder” class, one for each instance that the
2 template creator wants created whenever a template user creates an instance of the
3 template class. The following is an example of the TemplateBuilder class.

```
4 Class MSFT_TemplateBuilder
5 {
6     [key] string Name;
7     [key] string Template;
8     object Target;
9     string ControllingProperty;
10    string NamespaceProperty;
11    uint32 Order;
12};
```

11 In the above example, “Name” is a unique name for this object within the
12 context of this template. “Template” is the name of the template class for which
13 this template builder is specified. “Target” is an embedded object that will be
14 instantiated (PutInstance() will be called with it as an argument).
15 ControllingProperty (optional) makes instantiation of the builder be dependent on
16 a template argument. If the ControllingProperty property is specified, then on
17 instantiation, the template provider will observe the template argument named by
18 the property. If the value of this property is NULL or if the property is a boolean
19 type and its value is “false”, then the builder will not be instantiated.
20 NamespaceProperty (optional) is the name of a property on the template object.
21 At instantiation time, the value of the NamespaceProperty specifies the namespace
22 in which the target will be instantiated. “Order” (optional) identifies the order in
23 which the template provider instantiates TemplateBuilders. The Order value is an
24 integer that specifies that all builder objects having an Order X will be instantiated

1 before builder objects having an Order Y where X < Y. If no order is specified
2 then it is defaulted to 0.

3 If the instance to be created by the builder is not parameterized by any
4 template arguments, TemplateObject is simply the instance to be created, without
5 any special qualifiers. An example TemplateObject is:

```
6 Instance of MSFT_TemplateBuilder
7 {
8     Name = "ExampleBuilder";
9     Template = "CorrelationTemplate";
10    TemplateObject = instance of EventFilter
11    {
12        Query = "select * from InstanceModificationEvent WHERE
13 TargetInstance ISA 'MyClass'"
14    };
15    };
16
17
18
19
20
21
22
23
24
25
```

14 In most cases, the values in the object to be created are parameterized by the
15 values in the template. For instance, in the above example, the template creator
16 would likely want to register a filter for events of the type specified in the
17 WatchClass property of the template. Three instance-level property qualifiers are
18 available to allow parameterization of objects. These three instance-level property
19 qualifiers are identified below.

21 [tmpl_prop_val(PropertyName)] specifies that this property should
22 receive the value of the property PropName of the template class, as in
23 [tmpl_prop_val(SomeValue)]

1 [tmpl_subst_str(SubstitutionString)] specifies that this property
2 should receive the value of the SubstitutionString with all the substitutions
3 performed. SubstitutionString uses environment-variable-like syntax:

4
5 [tmpl_subst_str("select * from %WatchClass%")] SomeProperty;
6

7 Here, SomeProperty will receive "select * from MyClass" if the WatchClass
8 property of the template class had a value of 'MyClass'. If the property specified
9 by WatchClass is null, then no substitution takes place. In this case,
10 SomeProperty would be 'select * from'.

11 To support more complex string substitutions, extension functions can be
12 placed in the substitution string. More complex substitutions arise when dealing
13 with substitutions involving SQL queries. Two extension functions supported by
14 the template provider are:

15
16 %!ConditionalSubstitution("String", TemplatePropIdentifier)%

17 and

18 %!PrefixedWhereSubstitution(Identifier, QueryTemplatePropIdentifier)%

19
20 The ConditionalSubstitution extension function allows "String" to be
21 substituted if the template instance property specified by Identifier is not null.
22 This is useful for templates that form SQL queries which may or may not have a
23 where clause. The conditional substitution in this example would be the
24 'WHERE' keyword. Using the original example, this might look like ...

1 [tmpl_subst_str(“select * from %WatchClass% “
2 “%!ConditionalSubstitution(“WHERE”, WatchScopeExpr)% “
3 “%WatchScopeExpr% “)] SomeProperty;

4
5 The reason why it is not sufficient to require the WatchScopeExpr parameter to
6 just contain the WHERE keyword, is that often, the where clause is also used in
7 queries which already have a where clause. For these cases, it would be necessary
8 to prepend an AND to the where clause. Using the ConditionalSubstitution
9 extension function solves these types of problems.

10 The PrefixWhereSubstitution extension function allows the property
11 references of a VALID ‘where’ clause of a SQL query to be prefixed with the
12 value of ‘Identifier.’ The TemplateQueryPropIdentifier refers to a template
13 parameter having a valid ‘where’ clause as its value. A valid where clause is the
14 substring of a valid SQL statement after the where keyword. An example of the
15 ‘PrefixWhere’ substitution would convert the clause “A = 1 and B=2 or C=3” to
16 “MyIdentifier.A=1 and MyIdentifier.B=2 or MyIdentifier.C=3”. If
17 ‘WatchScopeExpr’ in the preceding example template class had a value of “A=1
18 and B=2 or C=3”, then:

19
20 [tmpl_subst_str(“select * from InstanceModificationEvent “
21 “WHERE TargetInstance ISA ‘%WatchClass%’ “
22 “%!ConditionalSubstitution(“AND”, WatchScopeExp
23 “%!PrefixWhereSubstitution(TargetInstance, WatchScopeExpr)%”]
24 SomeProperty;

1 would set the value of SomeProperty to ...

2
3 select * from InstanceModificationEvent WHERE TargetInstance ISA
4 'MyClass' AND TargetInstance.A = 1 and TargetInstance.B=2 or
5 TargetInstance.C=3

6
7 When using both qualifiers, key properties of the target objects created by
8 other builders of the template can be referenced using the
9 BUILDER.<BuilderName>.<KeyName> syntax. RELPATH is a valid keyname.
10 This allows one to use key properties of instantiated targets that are 'keyholed'
11 when creating other target objects.

12 It should be noted that there is nothing preventing the TemplateObject
13 embedded object property of the TemplateBuilder from being yet another
14 Template. This allows templates to be nested within other templates.

15 16 Template Associations

17 A Template instance can be associated with the instances that the Template
18 Builder objects create on instantiation and vice-versa. This association is defined
19 as follows:

20
21 MSFT_TemplateAssociation
22 {
23 [key] MSFT_TemplateBase ref Template;
24 [key] object ref Target;
25 };

1 The template provider impersonates the client of its methods. This ensures
2 that actions that are performed by the template provider, such as instance creation,
3 will be done using the identity of the caller.

4 The template provider returns complex error information when creating
5 template instances. This information is provided in the form of a COM Error
6 object. The error object implements IWbemClassObject and its definition looks
7 like:

```
8        class TemplateErrorStatus : ExtendedStatus  
9        {  
10            string Property;  
11            string ErrorStr;  
12            object TemplateBuilder;  
13            object TemplateObject;  
14            object ExtendedStatus;  
15        };
```

16 The Property attribute names the property of the TemplateObject that the
17 Template Provider was trying to perform a substitution on before instantiation.
18 Can be NULL (when substitution has been performed correctly, but the
19 PutInstance() failed). The ErrorStr attribute can contain a hint as to what caused
20 the error. This property is used to contain extra error information (can be NULL).
21 TemplateBuilder is the builder object that the template provider was working on
22 when the error occurred. TemplateObject is the object that is instantiated by the
23 builder. All substitutions will have been performed at this point. Can be NULL
24 (when there is an error in substitution). The ExtendStatus field allows the error
25 object (if any) resulting from a failed instantiation of a TemplateBuilder. This
 error object is created by the provider that is the instance provider for the object

1 held in the TemplateBuilder. Can be NULL (if there is an error in substitution or
2 if the provider that backs the template object does not return an extended error
3 object). In the case of nested templates, the ExtendedStatus object could be
4 another instance of TemplateErrorStatus.

5

6 Correlation Templates

7 As discussed above, template functionality is independent of the updating
8 consumer functionality. There are, however, two important advantages with using
9 WMI Templates when constructing low-level correlation scenarios. These
10 advantages are:

11 **Parameterization** – Low-level correlation scenarios can be parameterized so
12 that they can apply across arbitrary input parameters. For example, one can
13 construct a ‘throughput’ template that can be applied to any numeric
14 property of any class.

15 **Manageability** - A single template instance can represent all the instances
16 that make up a low level correlation scenario (updating consumer instance,
17 event filters, bindings, etc.). This means that manipulation of this instance
18 affects all the instances of the scenario. For example, one can activate,
19 deactivate, or delete a template instance, thereby activating, deactivating, or
20 deleting all the instances associated with the low-level correlation scenario.

21 There are different ways in which templates can be used with WMI Event
22 Correlation, such as Correlation Helper Templates and Correlation Scenario
23 Templates.

1 Correlation Helper Templates

2 A template can be used to parameterize a particular aspect of a correlation
3 scenario. In this role, the template does not define a complete correlation scenario.
4 A distinguishing characteristic of this type of template is that it does not define
5 any correlation state. Although this type of template is referred to as a
6 “correlation helper template”, it can be used in situations other than event
7 correlation situations.

8 An example of a correlation helper template is one that handles the
9 instantiation of EventFilter and FilterToConsumerBinding instances given a
10 consumer ref and filter query string.

11 Correlation Scenario Templates

12 A correlation scenario template is one that encapsulates a specific
13 correlation scenario. Unlike a correlation helper template, a correlation scenario
14 template is autonomous. It maintains its own correlation state and is comprised of
15 correlation primitives, correlation helper templates, and even other correlation
16 scenario templates. A desirable feature of the correlation scenario templates is
17 that they can act as building blocks that a high-level user (e.g., a system
18 administrator) can assemble to create a larger, more complex, correlation scenario
19 template. In order to realize this goal, correlation scenario templates must be
20 designed with well-defined characteristics that allow them to be combined in
21 many, often unpredictable, ways.

Correlation Helper Template Design Guidelines

There are a few guidelines for developing helper correlation templates. These guidelines consist of a set of qualifiers that can aid tools that support instantiation of helper correlation templates. Since correlation template instantiation will mostly be instantiated by correlation scenario templates, these qualifiers (shown below in Table 2) will most likely be used for correlation scenario template design.

TABLE 2

Description	Class-Level	Description String
crln_data_query	Prop-Level	valid WQL data query
crln_update_query	Prop-Level	valid UpdateQL query
crln_event_query	Prop-Level	valid WQL event query
crln_namespace	Prop-Level	valid namespace
crln_classname	Prop-Level	valid class name
crln_condition	Prop-Level	valid WQL expression for constraints
crln_domain	Prop-Level	valid WQL expression for restricting the domain of an input
crln_propname	Prop-Level	valid property name
crln_delay_tol	Prop-Level	Interval expressed in floating point seconds
crln_interval	Prop-Level	Interval expressed in floating point seconds

1 All property level qualifiers can optionally contain a ‘grouping’ name as a value.
2 This value allows different aspects of a common entity to be related. For example,
3 a class name and property name may be related using a grouping identifier.

4

5 Domain and Condition Expressions

6 A domain expression is one that restricts the domain of the template input.
7 An example of a domain expression might be Name = ‘Foo’. This domain
8 expression tells the template to only consider instances having a scenario property
9 of ‘Foo’. A condition expression is one that identifies a particular condition of the
10 template input that is of interest to the template. An example of this might be a
11 template which performs actions on transition into a ‘True’ and ‘False’ state. The
12 condition describing the True state is specified by a condition expression. The
13 template can then automatically determine the False state by performing a
14 NOT(condition_expression). The distinction between a domain expression and a
15 condition expression allows efficient filtering to be performed by the correlation
16 template.

17 For example, imagine that there was only one expression passed to the
18 template described above.

19
20 (Name = ‘Foo’ OR Name = ‘Bar’) AND Prop1 > 50

21 Here, the domain expression would normally be Name = ‘Foo’ OR Name = ‘Bar’
22 and the condition expression would be Prop1 > 50. The implementation of the
23 template is interested when Prop1 goes below 50 as well so it could perform the
24

1 action that corresponds to the FALSE state. A template could do this through an
2 event filter expression such as:
3

```
3     Select * from InstanceModificationEvent where TargetInstance ISA  
4 'MyClass' AND NOT( %PropScopeExpr% )  
5
```

6 But, this means that the updating consumers subscribed to this event will be
7 indicated for instances that do not have a name 'Foo' or 'Bar'. The correct way to
8 issue this query would be to split the domain and condition expressions and form a
9 filter such as:
10

```
11     Select * from InstanceModificationEvent where TargetInstance ISA  
12 'MyClass' AND %ScopeExpr% AND NOT ( %PropExpr% )  
13
```

14 Correlation Helper Template Example

15 The following is an example of a correlation helper template:
16

```
17 //*****  
18 // ConsumerFilterBindingTemplate  
19 //*****  
20 /  
21 [  
22     dynamic,  
23     provider("Microsoft WMI Template Provider"),  
24     description("Creates a Binding and EventFilter given a Consumer and"  
25         "Event Query.")  
26 ]  
27 class ConsumerFilterBindingTemplate : MSFT_CorrelationHelperTemplate  
28 {  
29     [NotNull] EventConsumer ref Instruction;  
30     [crln_event_query, NotNull] string FilterQuery;  
31     boolean DeliverSynchronously;  
32 };  
33 instance of MSFT_TemplateBuilder
```

```
1  {
2      Name = "FilterCreation";
3      Template = "ConsumerFilterBindingTemplate";
4      Target = instance of EventFilter
5      {
6          /* Use the Id of this template ( which is unique among all
7          correlation templates, as the Id of the filter. ) */
8          [tmpl_prop_val("Id")] Name;
9
10         /* Use the value FilterQuery property to set the Query property
11         on the target filter object */
12         [tmpl_prop_val("FilterQuery")] Query;
13
14         QueryLanguage = "WQL";
15     };
16     Order = 1;
17 };
18
19 instance of MSFT_TemplateBuilder
20 {
21     Name = "BindingCreation";
22     Template = "ConsumerFilterBindingTemplate";
23     Target = instance of FilterToConsumerBinding
24     {
25         /* This refers to the relpath of the target instance created by
26         the Filter Creation builder. Note that when doing this, you must ensure
27         that this builder is instantiated after the FilterCreation builder.
28         This is done using the Order property on the builder. */
29
30         [tmpl_prop_val("BUILDER.FilterCreation.RELPATH")] Filter;
31
32         [tmpl_prop_val("Instruction")] Consumer;
33         [tmpl_prop_val("DeliverSynchronously")] DeliverSynchronously =
34         FALSE;
35     };
36     Order = 2;
37 };
38
39
```

Correlation Scenario Template Design Guidelines

Correlation scenario templates parameterize correlation scenarios. Thus, correlation scenario templates should follow all the guidelines for designing scenarios using correlation primitives. They should also follow the guidelines for designing helper correlation templates. Design guidelines for correlation scenario templates impose even more structure than helper correlation template design so that templates can be easily composed. There are two categories of correlation

1 scenario templates. These are ones that use events as their input (event-based) and
2 ones that use a data set as their input (data-based).

3

4 Event-Based Correlation Scenario Templates

5 Event-based correlation scenario templates allow for the correlation of
6 events. With event-driven scenarios, users are concerned with things such as how
7 often an event occurs, if the event has occurred within a certain amount of time,
8 etc. An important characteristic of event-driven correlation scenario templates is
9 that they typically do not need to be initialized. An event occurs and its action is
10 taken immediately (or soon thereafter). Typically, the user is unable to discover
11 the state of the inputs to this scenario, since events occur instantaneously and
12 cannot be reviewed at some later time.

13

14 Data-Based Correlation Scenario Templates

15 Although all correlation actions are taken as a result of some event (such as
16 a WMI event), the input to a certain class of correlation scenario templates is
17 better described by a data set description. Also, since data exists in the real world,
18 part of the correlation scenario template implementation is concerned with going
19 out and discovering the current state of the data. A data-driven correlation
20 scenario template is concerned with things such as whether A met a particular
21 condition for 5 minutes or what is the trend of A, etc. There are two subtypes of
22 data-based correlation scenario templates, property-based and condition-based.

Property-Based

Property-based correlation scenario templates are concerned with a specified property of the data set. This means that the scenario will maintain some state that is related to the value of the specified property. Examples of this type of scenario are trend or average correlation templates. These scenarios would be concerned with keeping the average or trend of a specified property. Property based correlation scenario templates usually require the property be of a certain type. For example the trend or average updating consumer scenario would require the property to be one of the numeric types.

Condition-Based

A condition-based correlation scenario template is concerned with watching a specified condition. This means that the scenario will keep some state that is based on the truth value of the condition over the specified data set. Examples of this type of scenario is seeing if a certain condition holds over each element of the data set for a specified amount of time, if a certain condition holds for a specified number of elements of the data set at any given time, etc.

Correlation scenario templates formally declare their input and output using class level qualifiers. The following table describes examples of these class level qualifiers.

TABLE 3

crln_arity	Class	Value is an integer specifying the number of inputs
crln_type	Class	Value must be DataProperty, DataCondition, or Event
crln_output_class	Class	Value is the name of the 'public' state of the scenario

1
2 The output data instances can be determined by using the class name of the output
3 class qualifier and constructing the following query:

4
5 Select * from OutputClass WHERE Scenario = 'ScenarioName'.
6

7 Correlation Scenario templates also define one or more property level
8 qualifiers that can appear on a template class property. This property is called the
9 correlation Id property. A Correlation Id is the unique id of the state instances of a
10 correlation scenario. This id is usually determined from the id of the data or the
11 events that the state instances represent. The correlation Id property is the name
12 of the property of the data or event that has a value that can be used as the
13 correlation Id. The name of this property, as input to a correlation scenario
14 template, helps the template implementation derive a unique id for the correlation
15 state instances. For correlation scenario templates having multiple instances per
16 scenario, this id will typically be the relative pathname of the data instance it
17 represents. Some scenarios have only a single instance per scenario. In these
18 cases, the correlation Id will most likely be the class name of the data instance(s) it
19 is representing.

20 An important point about correlation Ids is that they need to be the id of
21 the 'real-world' data instance it represents. In the cases where there is a single
22 correlation scenario template having state representing 'real-world' data instances,
23 this previous statement is obvious. However, when correlation scenario templates
24 have state that represents the data instances of other correlation scenario templates
25 (known as "layered correlation scenario templates"), the correlation Ids must be

1 that of the original ‘real-world’ data instances, not of the immediate correlation
2 data instances. In other words, the Id of the ‘real-world’ data instance must be
3 propagated when layering correlation scenario templates. This rule allows
4 correlation templates to be combined later using joining correlation scenario
5 templates.

6

7 Correlation Scenario Template Example

```
8 // Example classes and instances.  
9  
10 class ExampleClass  
11 {  
12     [key] string Name;  
13     boolean Prop;  
14 };  
15  
16 instance of ExampleClass  
17 {  
18     Name = "A";  
19     Prop = TRUE;  
20 };  
21  
22 instance of ExampleClass  
23 {  
24     Name = "B";  
25     Prop = FALSE;  
26 };  
27  
28 // the correlation scenario definition  
29  
30 [ dynamic, provider("Microsoft WMI Transient Provider") ]  
31 class ExampleCorrelationState : MSFT_CorrelationStateBase  
32 {  
33     boolean ReceivedEventA;  
34     boolean ReceivedEventB;  
35     [trns_egg_timer] uint32 Timer;  
36 };  
37  
38 Class BothAandBEvent : MSFT_UCEventBase  
39 {  
40     string Name;  
41 };  
42  
43 /* the template class */  
44  
45 [dynamic, provider("Microsoft WMI Template Provider")]  
46 class SimpleCorrelationTemplate : MSFT_CorrelationScenarioTemplate
```

```
1     {
2         [crln_event_query] string EventQueryA;
3         [crln_event_query] string EventQueryB;
4     };
5
6     /* the template builders */
7     instance of MSFT_TemplateBuilder
8     {
9         Name = "InitializerUC";
10        Template = "SimpleCorrelationTemplate";
11
12        Target = instance of MSFT_UpdatingConsumer
13        {
14            /* this naming scheme allows a user to determine the entire
15            ancestry of a template instance. The scope is the name of the parent
16            template. In our case, there will not be a parent template, but
17            templates should be designed with this in mind. Since the element
18            being created is an atomic instance (not a template) the scope and the
19            name fit into the name property. Note that this naming scheme is
20            optional and is purely for browsability. The real keys of target
21            instances will usually be keyholed. For instance, the key for an
22            updating consumer, the 'Id' parameter is keyholed here. This approach
23            keeps keys a reasonable size, which can be important, especially when
24            there are many nesting levels. */
25
26            [tmpl_subst_str("%Scope%!%CLASS%=%Name%!Initializer")] Name;
27
28
29            /* The Scenario should always be the Id of the creating
30            template. Following this rule allows a user creating the template to
31            know what the scenario property of the state instances will be. */
32
33            [tmpl_prop_val("Id")] Scenario;
34
35            /* When updating correlation state, again we use the Id of the
36            template as the scenario property of the state. This will allow
37            multiple instances of this template to be created with having their
38            state instances conflict with one another */
39
40            [tmpl_subst_str{ "INSERT INTO ExampleCorrelationState "
41                "( Id, Scenario, "
42                "ReceivedEventA, ReceivedEventB, Timer ) "
43                "( 'ExampleCorrelationTemplateState', "
44                "'%Id%', FALSE, FALSE, 0 ) "}] Commands;
45        };
46    };
47
48    instance of MSFT_TemplateBuilder
49    {
50        Name = "EventAUC";
51        Template = "SimpleCorrelationTemplate";
52
53        Target = instance of MSFT_UpdatingConsumer
54        {
55            [tmpl_subst_str("%Scope%!%CLASS%=%Name%!SetEventA")] Name;
56            [tmpl_prop_val("Id")] Scenario;
```



```
1 instance of SimpleCorrelationTemplate
2 {
3     Id = "ExampleCorrelationTemplate";
4     Name = "ExampleCorrelationTemplate";
5
6     EventQueryA = "SELECT * FROM InstanceModificationEvent "
7         "WHERE TargetInstance ISA 'ExampleClass' "
8         "AND TargetInstance.Name = 'A' ";
9
10    EventQueryB = "SELECT * FROM InstanceModificationEvent "
11        "WHERE TargetInstance ISA 'ExampleClass' "
12        "AND TargetInstance.Name = 'B' ";
13
14};

15
16
17
18
19
20
21
22
23
24
25
```

Fig. 9 is a flow diagram illustrating a procedure 900 for combining two templates to create a third template. Initially, a template creator creates a first template class for handling events (block 902). The template creator then designates inputs associated with the first template class (block 904) and designates outputs associated with the first template class (block 906). The template creator also creates a second template class for handling events (block 908). The template creator then designates inputs associated with the second template class (block 910) and designates outputs associated with the second template class (block 912). The template creator creates a third template class by combining the first template class and the second template class (block 914).

An example template class and instances of the template class are provided below.

```

1  ****
2  *
3  *      ConsumerFilterBindingTemplate
4  ****
5  */
6
7  [
8  dynamic,
9  provider("Microsoft WMI Template Provider"),
10 description("Creates a Binding and EventFilter given a Consumer and
11 "
12         "Event Query.")
13 ]
14 class ConsumerFilterBindingTemplate : MSFT_CorrelationHelperTemplate
15 {
16     [NotNull] __EventConsumer ref Instruction;
17     [crln_event_query, NotNull] string FilterQuery;
18     boolean DeliverSynchronously;
19 };
20
21 instance of MSFT_TemplateBuilder
22 {
23     Name = "FilterCreation";
24     Template = "ConsumerFilterBindingTemplate";
25     Target = instance of __EventFilter
26     {
27         /* Use the Id of this template ( which is unique among all
28 correlation
29             templates, as the Id of the filter. ) */
30         [tmpl_prop_val("Id")] Name;
31
32         /* Use the value FilterQuery property to set the Query property
33 on
34             the target filter object */
35         [tmpl_prop_val("FilterQuery")] Query;
36
37         QueryLanguage = "WQL";
38     };
39     Order = 1;
40 };
41
42 instance of MSFT_TemplateBuilder
43 {
44     Name = "BindingCreation";
45     Template = "ConsumerFilterBindingTemplate";
46     Target = instance of __FilterToConsumerBinding
47     {
48         /* This refers to the relpath of the target instance created by
49 the
50             Filter Creation builder. Note that when doing this, you
51 must ensure
52             that this builder is instantiated after the FilterCreation
53 builder.
54             This is done using the Order property on the builder. */
55
56         [tmpl_prop_val("__BUILDER.FilterCreation.__RELPATH")] Filter;
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75

```

```
1      [tmpl_prop_val("Instruction")] Consumer;
2      [tmpl_prop_val("DeliverSynchronously")] DeliverSynchronously =
3      FALSE;
4      };
5      Order = 2;
6  };
7
8
9
10
11
12
13
14
15
```

Fig. 10 illustrates an example of a suitable operating environment in which the template creation system and method may be implemented. The illustrated operating environment is only one example of a suitable operating environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Other well-known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, gaming consoles, cellular telephones, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

Fig. 10 shows a general example of a computer 1000 that can be used in accordance with the invention. Computer 1000 is shown as an example of a computer that can perform the various functions described herein. Computer 1000 includes one or more processors or processing units 1002, a system memory 1004, and a bus 1006 that couples various system components including the system memory 1004 to processors 1002.

The bus 1006 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. The system memory 1004 includes read only memory (ROM)

1 1008 and random access memory (RAM) 1010. A basic input/output system
2 (BIOS) 1012, containing the basic routines that help to transfer information
3 between elements within computer 1000, such as during start-up, is stored in ROM
4 1008. Computer 1000 further includes a hard disk drive 1014 for reading from
5 and writing to a hard disk, not shown, connected to bus 1006 via a hard disk drive
6 interface 1015 (e.g., a SCSI, ATA, or other type of interface); a magnetic disk
7 drive 1016 for reading from and writing to a removable magnetic disk 1018,
8 connected to bus 1006 via a magnetic disk drive interface 1019; and an optical
9 disk drive 1020 for reading from and/or writing to a removable optical disk 1022
10 such as a CD ROM, DVD, or other optical media, connected to bus 1006 via an
11 optical drive interface 1023. The drives and their associated computer-readable
12 media provide nonvolatile storage of computer readable instructions, data
13 structures, program modules and other data for computer 1000. Although the
14 exemplary environment described herein employs a hard disk, a removable
15 magnetic disk 1018 and a removable optical disk 1022, it will be appreciated by
16 those skilled in the art that other types of computer readable media which can store
17 data that is accessible by a computer, such as magnetic cassettes, flash memory
18 cards, random access memories (RAMs), read only memories (ROM), and the
19 like, may also be used in the exemplary operating environment.

20 A number of program modules may be stored on the hard disk, magnetic
21 disk 1018, optical disk 1022, ROM 1008, or RAM 1010, including an operating
22 system 1028, one or more application programs 1030, other program modules
23 1032, and program data 1034. A user may enter commands and information into
24 computer 1000 through input devices such as keyboard 1036 and pointing device
25 1038. Other input devices (not shown) may include a microphone, joystick, game

1 pad, satellite dish, scanner, or the like. These and other input devices are
2 connected to the processing unit 1002 through an interface 1026 that is coupled to
3 the system bus (e.g., a serial port interface, a parallel port interface, a universal
4 serial bus (USB) interface, etc.). A monitor 1042 or other type of display device is
5 also connected to the system bus 1006 via an interface, such as a video adapter
6 1044. In addition to the monitor, personal computers typically include other
7 peripheral output devices (not shown) such as speakers and printers.

8 Computer 1000 operates in a networked environment using logical
9 connections to one or more remote computers, such as a remote computer 1046.
10 The remote computer 1046 may be another personal computer, a server, a router, a
11 network PC, a peer device or other common network node, and typically includes
12 many or all of the elements described above relative to computer 1000, although
13 only a memory storage device 1048 has been illustrated in Fig. 10. The logical
14 connections depicted in Fig. 10 include a local area network (LAN) 1050 and a
15 wide area network (WAN) 1052. Such networking environments are
16 commonplace in offices, enterprise-wide computer networks, intranets, and the
17 Internet. In certain embodiments, computer 1000 executes an Internet Web
18 browser program (which may optionally be integrated into the operating system
19 1028) such as the “Internet Explorer” Web browser manufactured and distributed
20 by Microsoft Corporation of Redmond, Washington.

21 When used in a LAN networking environment, computer 1000 is connected
22 to the local network 1050 through a network interface or adapter 1054. When
23 used in a WAN networking environment, computer 1000 typically includes a
24 modem 1056 or other means for establishing communications over the wide area
25 network 1052, such as the Internet. The modem 1056, which may be internal or

1 external, is connected to the system bus 1006 via a serial port interface 1026. In a
2 networked environment, program modules depicted relative to the personal
3 computer 1000, or portions thereof, may be stored in the remote memory storage
4 device. It will be appreciated that the network connections shown are exemplary
5 and other means of establishing a communications link between the computers
6 may be used.

7 Computer 1000 typically includes at least some form of computer readable
8 media. Computer readable media can be any available media that can be accessed
9 by computer 1000. By way of example, and not limitation, computer readable
10 media may comprise computer storage media and communication media.
11 Computer storage media includes volatile and nonvolatile, removable and non-
12 removable media implemented in any method or technology for storage of
13 information such as computer readable instructions, data structures, program
14 modules or other data. Computer storage media includes, but is not limited to,
15 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
16 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
17 tape, magnetic disk storage or other magnetic storage devices, or any other media
18 which can be used to store the desired information and which can be accessed by
19 computer 1000. Communication media typically embodies computer readable
20 instructions, data structures, program modules or other data in a modulated data
21 signal such as a carrier wave or other transport mechanism and includes any
22 information delivery media. The term "modulated data signal" means a signal that
23 has one or more of its characteristics set or changed in such a manner as to encode
24 information in the signal. By way of example, and not limitation, communication
25 media includes wired media such as wired network or direct-wired connection,

1 and wireless media such as acoustic, RF, infrared and other wireless media.
2 Combinations of any of the above should also be included within the scope of
3 computer readable media.

4 The invention has been described in part in the general context of
5 computer-executable instructions, such as program modules, executed by one or
6 more computers or other devices. Generally, program modules include routines,
7 programs, objects, components, data structures, etc. that perform particular tasks
8 or implement particular abstract data types. Typically the functionality of the
9 program modules may be combined or distributed as desired in various
10 embodiments.

11 For purposes of illustration, programs and other executable program
12 components such as the operating system are illustrated herein as discrete blocks,
13 although it is recognized that such programs and components reside at various
14 times in different storage components of the computer, and are executed by the
15 data processor(s) of the computer.

16 Although the description above uses language that is specific to structural
17 features and/or methodological acts, it is to be understood that the invention
18 defined in the appended claims is not limited to the specific features or acts
19 described. Rather, the specific features and acts are disclosed as exemplary forms
20 of implementing the invention.